

Educom Review

42 It's the Student, Stupid!

Vinton Cerf

Poet-Philosopher of the Net

36

44 So Wide a Web (So Little Time)

24 Why Technology? Provocative Answers from Six Leaders

By Laurence Alvarez • Carole Barone • Polley Ann McClure • Martin Ringle • John Stuckey • Thomas W. West

**Don Tapscott's Themes
for New Learning** 52

A Platonic Paradox 60

**Free Speech
& the Public Academy** 22

**The End of the Sixties
& the Future
of Technology** 6


EXXplicit Material Online

**Robert O'Neil's
Views on Free Expression
in the Information Age**

16

May
\$3.95


KATIE HAFNER
NEWSDEK
101 N 6TH ST SUITE 802
AUSTIN TX 78701-2932



VINTON Cerf is senior vice president of data architecture for MCI Engineering. Cerf, who was with the company in the early '80s, currently is responsible for developing the network architecture of MCI's future data services.

Prior to rejoining MCI, Cerf served as vice president of the Corporation for National Research Initiatives (CNRI) where he conducted national research efforts on information infrastructure technologies. Cerf co-developed the computer networking protocol, TCP/IP, widely used in the industry and for communications between the diverse university, government and commercial data networks, known collectively as the Internet.

From 1982 to 1986, Cerf was vice president of MCI Digital Information Services serving as chief engineer of MCI Mail. He also played a major role in sponsoring the development of Internet-related data packet technologies during his stint with the Department of Defense's Advanced Research Projects Agency (ARPA). He served on the Stanford University faculty in the computer science and electrical engineering departments and worked on the ARPA network development at UCLA.



poet- philosopher of the Net

EducOm Review: Let's start off with a poetical question, because George Gilder, we noticed, recently referred to you in *Forbes ASAP* as the poet-philosopher of the Internet.

Cerf: I didn't know that.

ER: That's a very nice compliment, and it inspires us to ask: when you were creating TCP/IP as the groundwork of the Internet, about 22 years ago, what did you think you were doing? Did you realize that you were building a cathedral? Or did you think you were just doing masonry?

Cerf: It's interesting that you should ask this particular question, because I'm right in the middle of reading through the manuscript of Katie Hafner's book on the Internet, and there's a wonderful quote from Paul Baran, who describes what he was thinking when he was doing his own work about 10 years earlier, in '62. He points out that when you're building a cathedral you can be easily persuaded that your piece of the cathedral is the most important part as you're working on it—yet when you step back and see what's going on, it's clear to you that nearly everybody had a block here or a block there, and a corner someplace else, but that it adds up in the end to a cathedral.

So really at the time that I was doing this work with Bob Kahn, my view of it was just trying to solve a very particular problem. There were successful demonstrations of the ARPANET—in fact the most successful public one had just taken place in October '72. Bob Kahn had gone to ARPA and started working on packet radio and packet satellites, applications of the same packet switching idea in different media; and when he came out to Stanford in early '73, he raised the problem of

how to interconnect those various nets, and so I was completely and totally focused on this specific technical problem. The model I had of all this was not building a cathedral or building a gigantic global public service, but it was simply trying to find a way to let the military build different kinds of communications networks that would serve computers in the field and allow all the various computers on the different networks to interwork with each other pretty much seamlessly. It was very much a technical focus and I didn't really have any expectations—at the time anyway—that it would turn into what it has. In fact I expected that if it were commercialized at all—and I don't think I was thinking about that in '73—that it would be completely reimplemented. All we were doing was groundwork—research to understand how the technology would work.

ER: When did you decide that it was a cathedral?

Cerf: Well, I didn't wake up one morning and say, "Oh, my goodness, we built a cathedral." For me anyway, the turning point started to come in the mid-eighties, believe it or not, and it was only after we started to see companies like Cisco Systems and Proteon and others start to make commercial products out of the routers—what we used to call gateways—that I began to realize that maybe we weren't ever going to get a chance to redesign this thing: it was just going to turn into a system that would grow without a complete revision. So I don't think I foresaw that the thing we had designed, and that built itself, would turn into what it is until about '86, and I think even then it wasn't very clear until about two or maybe four years later—about 1990. If you look at the growth curve of the Internet it starts to take off exponentially around 1988, although you don't often see until later



**"LEGISLATIVE TOOLS ARE VERY OFTEN NOT VERY DELICATE
INSTRUMENTS. THEY ARE GREAT BLUDGEONS,
AND THEY CAN INADVERTENTLY HAVE
ALL KINDS OF TERRIBLE SIDE EFFECTS."**

that it is an exponential buildup.

It turns out that Dan Lynch started this Interop Conference around '86, and I went to the first one and gave a talk there. There were maybe 300 people there. Not a very big thing. But the next year it was 900 people and so that meant in 1987 we started to think, "Gee, you know, a lot of people are interested in this." But even that didn't really do it. The first time I stopped dead in my tracks and understood that we were on to something really big was around '89, when Lynch's Interop attracted something like 15,000 people and had huge displays from all the various computer and router vendors. The growth curve went something like 300—900—5,000—15,000—and then some insane number like 35,000. And there is no question in my mind that it was walking into one of the Interops at which there were 15,000 attendees that caused me to just stop dead and say, "Oh my God, this is real."

ER: Other than the immense size and grandeur of the Internet, does anything else surprise you—or disappoint you?

Cerf: You mean about this whole evolution?

ER: Yes.

Cerf: I guess the surprises come because the spread of the system has been so rapid in unexpected quarters. And so when I get e-mail from people who are in obscure places around the world saying, "Hi, I'm on the Net," I'm always surprised. When I get e-mail from somebody in China and somebody in Australia and somebody in Africa, I'm always stunned. I say, "My God, it got there too?" When old friends that I haven't seen—elementary school or junior high school classmates like one that just popped up a couple of weeks ago—I got an e-mail from a guy in Guam I was in junior high school with, saying, "Hi, I'm teaching at the University of Guam and I'm on the Internet and how the heck are you?"—I always get blown away. Those are the kinds of surprises that are happening now, and of course the new applications that have come along, like Internet Multimedia and World Chat, are surprises just because I almost invariably underestimate the amount of human creativity there is out there, especially with millions of people trying things out. The disappointments come when I read things like the legislative efforts to restrict content on the Internet.

ER: You mean such as pornography?

Cerf: And it's not that I am in favor of pornography. That's not the issue. It is just that legislative tools are very often not very delicate instruments. They are great bludgeons, and they can inadvertently have all kinds of terrible side effects that destroy the utility of the medium. We've had these arguments for a hundred, two hundred years. When we talk about censorship in a new medium,

we almost always conclude that censorship is worse than the alternative. And so I am disappointed with some of the predictable reactions to this kind of an infrastructure. In the long term I am optimistic because I think that, on an international scale, it is not something that you can legislate domestically. And I also believe that the utility will be so overwhelmingly essential over time that these kinds of constraints will ultimately have to be rethought. I don't think that it is wrong to try to find ways of limiting access at the periphery. That is to say, if a parent wants to have software that somehow restricts the scope of action of the youngsters, that's a reasonable objective, although I must say it is technically challenging to effect. And I am also reminded by kids in their teens that this problem arises in all media and it has never stopped many of them from getting what they are interested in if they want to get it. So trying to find ways of *a priori* restricting content and access, I think, is not going to be 100 percent effective. But I certainly would hate to see the vast potential for this kind of technology stopped inadvertently by mistaken legislation.

ER: What about purists of a different kind? There have been numerous old hands on the Internet who have been appalled at the commercialization of the Internet.


Cerf: I have never had that view. In fact I am probably counterculture in that regard, because I have always believed—even as far back as the 1980s—that the network technology could not spread and could not penetrate as far as I wanted to see it go unless it became commercialized, because I believed that the system had to be economically self-sustaining. I didn't think—and I don't think—it is something the U.S. government could sustain and pay for out of tax dollars, because it has such huge scope. It has to have the same scope as the telephone network, maybe greater scope in a sense, because there are a lot more things that will be on the network than just telephones. There will be light bulbs on the network, for pete's sake—the scope is so huge for it to realize its potential that it just has to be economically self-sustaining.

So I welcome ways of turning this into a commercially supportable service. To the people who complain that somehow or other this means that the cost will go up and usefulness will go down, I would say that, to the contrary, competition will drive prices down—to say nothing of the fact that there is a certain economy of scale here and that the more systems and users there are out there, the farther the cost will be spread across many varying services. As a result the financial load will be lighter on everyone.

ER: What do you think is the current status of security on the Internet? In general, and also with regard to financial transactions.

Cerf: Security in general varies dramatically across the net-

**"I HAVE ALWAYS BELIEVED . . . THAT THE
NETWORK TECHNOLOGY COULD NOT SPREAD . . .
AS FAR AS I WANTED TO SEE IT GO
UNLESS IT BECAME COMMERCIALIZED."**



work. The most critical aspect is the security of the server machines. But those are not the responsibility of the network so much as they are responsibility of the people who operate the service systems—and you will find variations from very well-organized, well-structured, well-protected machines to things that are just sitting there with their innards hanging out. And so my first disappointment is from people who say, "Well, the network should protect all of my computers." And the answer is, "No, no, you don't understand: the network can't protect your computers completely, and your computers will be open to invasion if you don't configure them properly." The network does have to protect itself, but separately from protecting the computers that are on the Net. The individual computers have to protect themselves.

And also we have varying degrees of protection. I think operators are becoming more and more aware of the need to secure the controls of the various parts of the Net. But just remember, there are 60,000 networks out there and doubling every year, so all of the operators have to come to the same conclusion—to carefully protect the individual pieces of the Net that they operate. On the host side, well-configured operating systems are critical, and perhaps, in the longer term, end-to-end security and application-level security will carry us the rest of the way.

To come back to your question about financial transactions, that's how financial transactions are being protected today, on an end-to-end basis at the application level. And I think that—with the proper application of cryptography and the proper maintenance of key management discipline—you can protect the financial transactions quite well. The part that is difficult is that governments—many governments, not just ours—restrict the export of high-grade cryptography, some of which is needed to protect these financial transactions. And there is a lot of policy debate over how to balance intelligence and law enforcement interests against the need for protection, not just for financial transactions but for business in general.

ER: Speaking of government—do you have any thoughts about the leadership of the two most prominent politicians in the area of information technology, Vice President Gore and Speaker Gingrich?

Cerf: Well, let's see. First of all, I have been rather impressed by the Vice President's awareness of technology and what it can do. Occasionally, I've sat in on hearings where he participated prior to his ascendancy as Vice President, when he was still Senator. It seemed to me he was often the best-prepared person in the room, including the experts who were testifying. I think that he senses a great deal of the potential of this kind of communications facility, and I think that contributed to President Clinton's recent announcement in San Francisco regarding his objective to get every classroom in the country wired up on the Internet by the year

2000. Now, whether that's possible to do, you know, remains open to some debate, but industry is now being challenged to find a way to do that at some reasonable cost. So I think that we owe a lot to that kind of leadership.

I don't know the Speaker's views quite so well. I know he is influenced by Toffler and some of the other Futurist community. He also seems to feel very positively about this kind of technology and its utility. And perhaps part of that comes from his background as a historian, in the sense that communications is a way of contributing to people's knowledge of what's going on, of understanding history. I think he appreciates the usefulness of the technology, and I know the Republican Party seems to be moving pretty quickly to adopt use of the Internet as a tool for communicating with the constituents. This isn't to say the Democrats aren't doing the same thing, but it seems like the Republicans are a little better organized at it. I think the Speaker is a friend of this kind of technology.

The only place where I am a little concerned is that Congress in general—the Republican-dominated Congress—has been less supportive of research funding from the government than I would like. There's a bogus argument about application-level research, or applied research, that gets tossed into the unwarranted argument that, "Oh, you are picking winners and losers, you shouldn't get involved in any kind of applied research." I don't buy that at all. I cite the ARPANET and the Internet as prime examples of the need to get all the way down to the point where you have a working system before you can reasonably expect industry to make the investment to create products and services, especially if you are going into an area where you don't know if there is a market and the technology is unproven.

Industry, except for the small start-up, tends to be very reluctant to get into major investments if they don't have a reasonably predictable payoff. And that of course is a terrible observation on the dominance of Wall Street thinking where quarterly profits are almost everything—except for craziness. The potential energy of companies like Netscape tends to drive the gold rush phenomenon on Wall Street for awhile, but then once you go public you can't ride on the potential energy anymore—you have to suddenly start building quarterly profits. And you get these two different evaluation metrics. First comes the gold rush—"My God, this is the best thing since sliced bread!" But then after you are public, a different bunch of bean counters shows up saying, "Tell me what your share is going to be a year from now." So Wall Street is a fickle friend in this business.

ER: Well, how do you think that specific gold rush will play out?

Cerf: Well, if you remember your history of gold rushes, it is often the case that the people who made their money in the gold



"STUDENTS WILL BE CONFRONTED WITH SUCH A MIXTURE OF MATERIAL OUT THERE THAT THEY CANNOT AND SHOULD NOT ASSUME THAT ANYTHING THEY GET OFF THE NET IS NECESSARILY CORRECT."

rush weren't the guys that were panning for gold, it was the guys that were selling the pans and the mules and the food and the tents and the shovels. So looking at it from the standpoint of an important Internet service provider, with my MCI hat on, I believe that the fellows in the communications companies stand to make a great deal out of the growth of the Internet, not because of a crazy gold rush phenomenon like you see in Wall Street, but because ultimately people are going to be relying on this on a day-to-day basis.

ER: Let's talk a little now about content. Do you see the Internet as currently content rich?

Cerf: It seems to me very clear that the Internet as a medium is potentially capable of absorbing telephone, radio and television functionality as time goes on and it gets bigger and it gets faster and the switches can carry more traffic. Technically I believe there could be a convergence of that sort. But let's set that aside for a moment and talk about content as such. A lot of people think mostly in terms of content on the World Wide Web or reachable through the World Wide Web. That is certainly an important part but by no means the only content that is on the system. There are millions of computers out there, some of which could be available for supplying computer services quite independent of gathering information. And there are of course files sitting out there in archives containing software and everything else that can be distributed. So I just want to be careful not to limit our thinking about content to facilities or capabilities to things like World Wide Web. The chat rooms and the news rooms and all these other capabilities are a part of the whole fabric of the Internet, to say nothing of an avalanche of e-mail.

But to come back to the principal window by which people see the content, mainly the Web, it really is quite fascinating. It is far less organized than a carefully constructed set of volumes like the *Encyclopaedia Britannica*. In fact I was reflecting on that this morning because I had to go down and read about the early history of the telegraph, and I thought for a moment to try an experiment, which is to go read the *Encyclopaedia Britannica* renditions, which are quite lively in the telecommunications section of the EB. It goes all the way back to the 1700s. And I thought, well, I wonder what would happen if I typed "telegraph" to a Web searcher and see what would pop up. My guess was that it would be a lot less well-organized, and of course that guess was correct.

But let's go on in that vein. I was in Fort Lauderdale at a Conference called TelEd95, and there were a couple of thousand teachers there, and one of the things we talked about was the fact that there was an enormous amount of content on the Internet but that it was not organized, not validated necessarily, and not edited. A lot of it had inaccuracies. Some of it was extraordinarily good

and well-organized, but some of it was very poorly organized and not corroborated. And we talked about that for awhile and realized that one of the things that this forces us to do—in addition to solving the routine but difficult problem organizing material like a librarian has to do in a library—is to teach students critical thinking. Students will be confronted with such a mixture of material out there that they cannot and should not assume that anything they get off the Net is necessarily correct, without doing additional work to verify that. And we need to be sure that they understand that very important skill of critical evaluation.

And of course the Net is not the only source of information that will confront them. Television, radio, books, magazines and the like are all sources of information and they all require critical analysis. If, as a result of being confronted with this mixture of uneven content on the Net, we wind up cranking out people who think more about what they read and hear and see, we will have very generally benefited our culture and others around the world.

ER: Who do you see doing that?

Cerf: Well, actually I think the same scenarios that have worked in the other media will almost certainly work here. Think about television, film, radio; think publishing, recordings, CD-ROM and so on—these are all ways of conveying information to consumers, and what we do is to look for high-quality content. How are you going to identify high-quality content? Well, sometimes it comes because we trust the source of the information. When I pick up the *Wall Street Journal*, I make an assumption that the editors, understanding what my information needs and interests are, will provide me with some useful and reasonably accurate data. Of course newspaper deadlines all contribute to the erosion of accuracy, but nonetheless I think I have a higher likelihood of reading something useful about the business world in the *Wall Street Journal* than I do in some tabloid. We find publishers of information establishing reputations as to the quality of the content. We are more likely to find accuracy in *Science* magazine, for example, than in, say, *Discover*. And I don't mean to impugn the guys at *Discover*. I just think that *Science* magazine is trying to produce scientifically reproducible content, while *Discover* is probably a little bit more likely to focus on something that is scientifically titillating. So I think we will find the same thing happening in the Internet. We will find organizations or individuals who act as sources of—or editors of—content, and we will learn to trust them.

Now you might raise the technical question: How do I know this isn't a spoof? How do I know somebody didn't just claim this came from *Science* magazine online? And probably the answer to that will tend to come, in extreme cases, by way of digital signatures and things of that sort. It's somewhat like putting the *Scien-*

"THE TYPICAL LAPTOP OR DESKTOP WILL ABSOLUTELY BE ABLE TO DELIVER VIDEO TO YOU."



tific American logo onto something, so that if someone produces a fake copy of the magazine, *Scientific American* will discover it and raise hell.

So it's just like television: you are likely to select commentators whose biases you believe are more accurate and more suited to your needs. You listen to certain radio programs; you might listen to NPR, for example, because you think it chooses information that is more accurate and more relevant to you than some other radio stations. Well, surely the same mechanisms that have worked in the other media will tend to work for the Internet.

Of course, the thing that's different about the Internet is that it is far easier for an individual voice to be heard than in any other medium we've ever had before, because it mixes together the mass distribution capability and individual private communications and everything in-between, and as a result we will have a larger diversity of information to choose from. And perhaps in that sense we will have more opportunity to take advantage of the deep thinking of individuals that we have normally never heard from. The downside is that we will also be subjected to more nut material than we ever had to deal with before.

ER: Do you see a convergence of the communications technologies? Do you see television being looked at on the same box that we are using for computers?

Cerf: Well, I actually have some mixed feelings about this. I don't exactly mean to reject the idea of convergence, but I am very skeptical of the apocalyptic view that people develop any time a new technology comes along—the view that the new technology will replace everything else. That's just not very common. I think what typically happens is that new technologies find niches where they are extremely effective; and the others, if they are at all useful, tend to survive anyhow—which is why we still have books and magazines and newspapers and television and radio and pre-recorded cassettes and so on and so forth. Granted, we don't have long-playing records so much anymore, because we have CDs, and that over time has really been a major replacement; but that's the exception to the general rule.


In the case of television, I think that the typical laptop or desktop will absolutely be able to deliver video to you and you may choose to use that medium; for example, in the middle of everything else, you might pull up the news report. You may actually search the network for certain kinds of content and have it delivered as video. However, I imagine that most of those desktop and laptop displays will not really be very satisfying, if what you want is entertainment. So you may very well turn to a much larger system of projection—CRT or flat panel or what-have-you—as a medium by which you enjoy some entertainment. I think it is not very likely that you will wind up using that big projection television set

as a workstation. Some MIT studies suggest that there is quite a big difference between the aspect ratios that you expect in a workstation compared to what you expect if you are ten feet away from a large projection unit. So I think you may find some drifting or crossing of the lines, so that you get some information access on a large projection television, with low resolution, large font, big buttons, and so forth—as opposed to reading and writing reports and pulling up high-resolution content.

ER: What timeline do you see for that? Ten years from now? Twenty? Thirty?

Cerf: Well, I don't have enough knowledge of the display technology to know what to say here. I can tell you that I do believe that the flat panel technology concepts will ultimately dominate even in the large screen format. In fact, I saw a small example of this recently: *Fortune* magazine did a seminar for about 50 CEOs and they set up about 50 workstations all connected to the Internet. And rather than use the large CRT displays, which were big 17-inch affairs, they used very nice flat panel displays, that took up a lot less room, generated less heat, and probably alleviated anyone's immediate worries about being CRT-irradiated and all that stuff. So I suspect that we will see large-screen high-resolution units in 10 or 20 years or even sooner. But what I'm not so sure about is whether it will be comfortable to use that kind of unit for text composition and e-mail reading—sitting in an easy chair or something like that 10 or 12 feet away.

Also there is a question of privacy. It is a little bit like whether to use the speaker phone or not, and people who use IT telephony with laptops that have speakers and microphones quickly realize that that's like being on speaker phone and is not very private. Of course, if you really wanted to use it as a telephone, there are cases where you might like to have a handset or a headset or something which is a little less of a broadcast unit.

Also, I have to argue that if the large entertainment set is expensive, as I expect it will be, you are going to find its use dominated by multiple simultaneous users—in other words, people who are enjoying a common entertainment experience or who are at Nintendo controls and are playing some kind of competitive game. It is less likely that such an expensive unit is going to be dedicated to the use of one person to do a particular task. So suppose Johnny wants to use it to do his homework and Dad wants to watch the six o'clock news and Mom needs it to do some reference work, what are you going to do—have everybody fighting each other over the 16-foot display in the entertainment room? No, of course not. You will have flat panel or other workstation displays scattered around the house for dedicated use by an individual. So I think that to expect that one system will be used for everything is nonsense, because that's just not practical. 

Part II of this interview with Vint Cerf will appear in the next issue.